The Immediate Reality of Accessible Robotics: A Unitree G1 Case Study on Consumer-Available, Modifiable Platforms and Their Dual-Use Implications

Independent Researcher. Berend Watchus. The Netherlands. July 1, 2025.

mailonlinebw@protonmail.com

Abstract

This **paper offers** a **foresight perspective** on the immediate challenges and opportunities presented by accessible robotics. Using the **commercially available Unitree G1 humanoid robot as a case study**, it explores **how readily available AI integrations**, such as Large Language Models and smartphone-based emotional detection, **enable advanced capabilities and dual-use scenarios that are already a present reality**. The **paper analyzes** the **inadequacy of current legal frameworks through a hypothetical lawsuit**, highlighting the **urgent need** for proactive **regulatory action** concerning **public safety, privacy, and the prevention of misuse**. It further emphasizes that the **responsible development of AI**, particularly through "architectures of meaning" and "synthetic insula" concepts, is crucial for **ensuring safe and ethically aligned human-robot interaction** and navigating the inevitable future of accessible robotics.

Keywords

Accessible Robotics, Humanoid Robots, Unitree G1, Dual-Use Technology, AI Integration, Large Language Models (LLMs), Ethical AI, Robot Regulation, Public Safety, Privacy, GDPR, EU AI Act, Foresight, Synthetic Insula, Architectures of Meaning, Immediate Challenges, Commercial Availability, Emotional Detection, Smartphone Integration, Advanced Capabilities, Hypothetical Lawsuit, Proactive Regulation, Regulatory Action, Misuse Prevention, Responsible AI, Ethical Alignment, Human-Robot Interaction, AI Control, AI Alignment, Non-Sentient AI, Autonomous Systems, Conscious Realism, Unified Model of Consciousness, Substrate-Agnostic, Neural Correlates of Consciousness, Explainable AI, Trustworthy AI, Civilian Robots, Industrial Automation, Government AI.

Foreword: Embracing Innovation with Prudent Foresight

The rapid advancement of robotics and artificial intelligence presents humanity with unprecedented opportunities. This **paper is rooted in a forward-looking**

perspective that champions technological progress while advocating for a proactive and intelligent approach to its implications. **It is not an alarmist or anti-tech stance, but rather a call for diligent foresight into the expanding uses of these powerful tools.** By rigorously examining the capabilities and accessible nature of technologies **like the Unitree G1, which serve as modifiable platforms for advanced AI**, we aim to foster a balanced dialogue that enables us to harness the immense potential for benevolent aid while responsibly mitigating unforeseen hazards through timely regulation and ethical consideration. **Understanding the full spectrum of a technology's application, from its intended benign uses to its potential for repurposing**, is essential for **shaping a future where innovation thrives safely and equitably**.

Unitree G1: A Case Study in Immediate Dual-Use Robotics

Scenario Premise: Yesterday's Tech, Today's Immediate Reality

This isn't a futuristic thought experiment; it's a present-day reality. Advanced humanoid robots like the Unitree G1 (approx. 35kg, €16,000–€20,000) are already commercially available to the general public via direct mail order, making them accessible at a consumer-friendly price point for a device of its complexity. Crucially, their inherent modifiability and their design as open platforms compatible for integrations and upgrades—leveraging readily available tools like smartphones and online Large Language Models (LLMs)—means that capabilities once confined to advanced labs are now within reach of enthusiasts and hobbyists.

This widespread accessibility and upgrade potential collapses the distinction between a "friendly helper" and a "bad humanoid robot" scenario. It's akin to combining the unpredictable nature of unchained zoo animals with the raw, untamed power of loose cannons, all controlled by the deceptively simple consumer-grade remote control interface. Just as a 3D printer can democratize industrial parts or repair toys but also be used for unregistered components, or a helpful drone for aid relief can be repurposed for surveillance or payload delivery, the Unitree G1 embodies this same immediate dual-use challenge. This test case stresses the urgent need for robust regulatory frameworks, safety protocols, and clear legal precedents now, not in some distant future. The absence of specific laws means current situations are handled reactively, often ambiguously, and without adequate preventative measures.

Part 1: The Technical Testing Scenario (Expanded Capabilities & Applications)

Robot: Unitree G1 Humanoid Robot (with standard perception sensors: cameras, depth sensors, IMU). Its open architecture and SDK allow for extensive software modifications and integrations, facilitating significant upgrade potential.

Owner/Operator: Alex (25 years old, private citizen, tech enthusiast, has purchased the G1 online and is proficient in programming custom integrations).

Location: A moderately busy public park pathway in a major EU city (e.g., Berlin, Paris, Rome, but unspecified for general EU applicability).

Objective: Alex wants to "test" the G1's significantly enhanced capabilities, demonstrating how off-the-shelf hardware, when combined with accessible software and cloud AI, can achieve advanced interactions far beyond its default programming.

Technical Integrations Implemented by Alex (Leveraging Readily Available Tech):

"Self-Recognition & Environmental Understanding" Module:

- **Implementation**: The G1's existing cameras feed real-time visual data to an on-board computer. This computer runs powerful, open-source computer vision algorithms (e.g., YOLO for object detection, image segmentation techniques)—readily available online. The processed visual data (e.g., "humanoid shape, matching self-model, located in reflective surface, 2 meters away") is then sent via standard API calls to a cloud-based Large Language Model (LLM) (e.g., OpenAI's GPT-4o or Google's Gemini Pro).
- **LLM Role:** The LLM is prompted to interpret the visual context semantically. It can confirm "Is this my reflection?" or "Is this another entity?", and answer higher-level questions like "What are the common objects in this park environment and how are people interacting with them?" This provides the robot with a more nuanced, "cognitive" understanding of its surroundings, a significant upgrade from its base capabilities.
- **Applications (Toward 'Greatness'):** This module allows the G1 to move beyond simple obstacle avoidance to genuinely understand its spatial

relationship to its environment, identify other people and their activities, distinguish between static objects and dynamic elements, and potentially navigate more "intelligently" based on high-level LLM interpretations of its surroundings and even social cues.

"Emotional Context & Empathetic Response" Module:

Implementation: Alex develops a companion smartphone application—built using commonly available SDKs—that is designed to detect the emotional state of nearby humans. This app utilizes the smartphone's microphone to perform real-time speech analysis, looking at cues such as vocal tone, speech rate, and pitch variation. This emotional data (e.g., "speaker A shows signs of distress," "speaker B sounds joyful, approaching robot cautiously") is then wirelessly transmitted from the smartphone to the G1's control system, where it is fed as additional context into the prompts sent to the central LLM.

- **LLM Role:** The LLM, now informed by both visual data from the robot and sophisticated emotional cues detected by the widely available smartphone app, generates contextually appropriate conversational responses and guides the G1's physical reactions. For instance, if the app detects a child nearby expressing fear, the LLM might guide the G1 to take a step back, emit a soft, pre-recorded, non-threatening "beep," and slightly lower its head, demonstrating an "empathetic" response as derived from the LLM's understanding of how to de-escalate fear, all without direct manual input from Alex for each specific action.
- **Applications (Toward 'Greatness'):** This integration enables the G1 to engage in more "emotionally intelligent" interactions, moving beyond simple programmed responses to provide highly nuanced and seemingly empathetic engagement. This dramatically expands its potential for social robotics, public assistance roles, or creating a more 'aware' and responsive public presence.

The Test Operation:

Alex activates the G1. Using a standard RC remote, Alex guides the G1 along the park path. The G1, however, isn't simply reacting to stick commands; its semi-autonomous nature and custom, readily integrated modules are constantly at play, demonstrating its significant upgrade potential:

Self-Recognition in Action: As the G1 passes reflective surfaces (e.g., a large glass building nearby), its "self-recognition" module actively engages with the LLM. The robot might briefly pause and "examine" its reflection, perhaps tilting its head in a human-like gesture based on LLM-derived "curiosity" behaviors—a complex interaction initiated by accessible AI and readily integrated software.

Environmental Awareness: When a group of people gathers on the path, the G1's LLM-enhanced understanding of "crowd" or "gathering" leads it to autonomously slow its pace, perhaps even taking a slightly wider berth than Alex explicitly commanded via the RC, relying on the LLM's interpretation of "safe navigation around dynamic obstacles" in a highly unstructured environment.

Emotional Responsiveness: A small child, startled by the robot's appearance, lets out a gasp. Alex's smartphone app detects the child's "fear" vocalization. This information is instantly fed to the LLM. The G1, instead of continuing its programmed path, autonomously halts, emits a soft, pre-recorded, non-threatening "beep," and slightly lowers its head, demonstrating an "empathetic" response as derived from the LLM's understanding of how to de-escalate fear, all without Alex having to manually execute these precise actions. This showcases how a robot can gain "emotional intelligence" through accessible, off-the-shelf integrations.

Part 2: The Law-Making and Safety Protocols Test Case (The Inadequacy of Current Law)

Scenario Continuation:

Despite Alex's intent for a benign test, the G1's presence causes a stir. While no one is directly harmed, multiple calls are made to the police. People express unease, some parents shield their children, and one elderly person complains about the "menace" obstructing the path. A police officer arrives, observes the semi-autonomous behavior and the robot's "responsive" actions, and, after repeated warnings to Alex to cease and remove the robot (which Alex tries to comply with but the G1's complex movements make immediate, silent cessation difficult), decides to confiscate the G1, citing general public safety and nuisance concerns.

The Lawsuit:

Alex sues to retrieve the robot. Alex's lawyer argues:

- "My client broke no specific law. There is no statute in [EU Member State] prohibiting a private citizen from operating a semi-autonomous humanoid robot on a public path under RC control, even with advanced AI integrations and upgrades."
- "No incident occurred, no one was physically harmed, and no property was damaged. The police action was disproportionate and based on unfounded 'general safety concern' without a clear legal basis."
- "The robot is CE marked; it's a legal product. Its operation, even with AI and external sensor feeds, is still ultimately controlled by my client via remote, making it akin to a large, sophisticated RC toy."
- "The manufacturer's warnings are advisory, not law. Adults should be free to use their property, especially when enhancing it with widely available software and apps."

Judge’s Reasoning (Jurisprudence of Decisional Law):

A judge, confronted with this novel case, would have to interpret existing broader legal principles in the absence of specific robot legislation. The judge would likely rule against Alex and uphold the confiscation, based on the following:

Police Powers to Maintain Public Order and Safety:

- Preventative Action: The judge would acknowledge the police's inherent power to act preventatively where there is a clear and foreseeable risk to public safety or a disturbance of public order, even without an immediate incident. The robot's weight (35kg), its complex semi-autonomous nature, its LLM-driven "interpretations" (from accessible online services), and the public's negative reaction (complaints, fear) would be strong evidence of such a risk and disturbance.
- Proportionality: The police gave warnings. Confiscation would be deemed a proportional response when warnings were insufficient to cease the activity deemed a public safety threat.

Product Misuse and Operator's Duty of Care:

Intended Use vs. Actual Use: The G1 is a commercial/research product. The manufacturer's disclaimers (which are legally significant for liability) clearly indicate its intended use. Operating it in a crowded public park, despite the CE mark for general safety, falls outside its intended use and creates an unreasonable risk for the public.

Negligence Standard: Even if no specific robot law exists, the judge would apply the "reasonable person" standard. A reasonable and prudent person would not operate a 35kg, semi-autonomous robot with unpredictable, LLM-driven behaviors and external sensor integrations (even via a smartphone app) in a shared public space without specific authorization, training, or safety protocols.

Emerging EU AI Act and The Unpreparedness of Current Law:

High-Risk AI (Current Interpretation): The judge would lean on the spirit and upcoming provisions of the EU AI Act. The G1, especially with its advanced perception, LLM integration, and external smartphone sensor data in a public setting, would undoubtedly be classified as a "high-risk AI system" under the Act. While full deployer obligations might be phasing in (expected to apply from mid-2025 onwards, with some provisions later), the Act's existence signifies the EU's clear intent to regulate such systems stringently due to their potential for harm.

Lack of Compliance: Alex, as a private individual, would not have met the rigorous requirements for deploying a "high-risk AI system" in public (e.g., adequate risk assessment and mitigation systems, high-quality datasets, human oversight measures, robustness, cybersecurity, and accuracy, or potential registration). The judge would view this as operating a high-risk system without due diligence, even if the explicit law isn't yet fully applicable.

Privacy Concerns (GDPR):

The G1's cameras and particularly Alex's smartphone app (collecting voice/emotional data from the public) would constitute processing of personal data.

The collection of emotion detection data from passers-by, especially via a smartphone and LLM, could constitute biometric and potentially health-related data, falling under Article 9 of the GDPR (special category processing), which demands stricter conditions for processing. The judge could also factor in potential GDPR violations, as this data would be collected without explicit consent or a clear legal basis from the public, especially vulnerable individuals.

Conclusion of the Lawsuit: The judge would conclude that Alex's operation of the Unitree G1 robot constituted an unreasonable danger and public nuisance due to its weight, semi-autonomous nature, its complex LLM-driven behaviors (even when remote-controlled), the unpredictable public environment, and the lack of specific safety regulations or permits for such activity. The confiscation would be upheld as a necessary measure to protect the public. Alex would not get the robot back without demonstrating a plan for safe, compliant operation in a controlled environment, likely after meeting future regulatory requirements.

Part 3: Enhanced Applications vs. Inevitable Risks

The Immediate Reality: Greatness and Hazard, Both Within Public Reach

This scenario is not science fiction. The G1 and similar robots are being bought by hobbyists, researchers, and small businesses today, often via simple mail order. The ease of integration and upgradeability using universally available tools like smartphones and online LLMs means that highly advanced capabilities are within reach of any tech-savvy individual. This widespread accessibility, combined with the inherent power of these machines, creates a dangerous regulatory vacuum. We cannot wait for severe accidents or widespread misuse before defining the rules for this already present technology.

From Helper to Hazard: The Dual-Use Trajectory, Driven by Modifiability:

Applications Toward Greatness (The "Friendly Helper"):

Personal Assistance & Social Support: With self-recognition and sophisticated emotional intelligence, robots like the G1 could offer highly personalized and empathetic companionship for individuals, potentially even in home healthcare or

senior living, providing direct support and combating loneliness.

Public Guidance & Education: Stationed in public spaces, a G1 could provide nuanced, context-aware information, guide complex navigation in crowded areas, or act as an interactive educational exhibit, its LLM helping it understand complex queries and its enhanced perception enabling better spatial awareness in dynamic environments.

Democratizing Robotics: Accessible platforms like the G1 accelerate research, development, and innovation in AI and robotics, allowing a wider community to contribute to solutions in fields from logistics to entertainment, much like 3D printers democratized manufacturing by allowing 'home printing' of complex parts.

Humanitarian & Agricultural Aid: Similar to drones providing aerial mapping for disaster relief or precise crop monitoring in agriculture, advanced ground robots could reach difficult terrains for search and rescue, deliver aid to isolated communities, or automate hazardous agricultural tasks.

Risks of Repurposing (Inevitable Accidents and Deliberate Harm):

Complex Accidents (Beyond Simple Malfunction): The greater the autonomy and "intelligence" derived from LLM integrations and external sensors, the more complex and less predictable an accident becomes, ranging from serious home and traffic incidents to severe workspace accidents and even unintentional human rights violations. An LLM might misinterpret a subtle non-verbal cue (e.g., a child's excited shriek vs. genuine distress) leading the robot to move or respond in an unintended or startling way, causing a fall or collision that is not a mechanical failure but a 'cognitive' misstep. These errors, driven by complex, software-defined behaviors, are harder to diagnose and prevent than simple hardware faults.

Unintended Privacy Intrusions: The continuous collection of visual, auditory, and emotional data from the public by robots with advanced sensors and LLM processing, often via personal devices like smartphones, creates immense privacy risks. Without strict regulation, this data could be misused, exposed, or even weaponized through mass surveillance.

Deliberate Misuse / Weaponization: A general-purpose humanoid robot, readily available to the public and capable of advanced mobility, strength, and sophisticated AI (including environmental understanding, self-awareness, and emotional recognition), represents a critical dual-use dilemma. As a foundational principle: anything a human with two arms and two legs can do, these humanoids could eventually replicate, whether intentionally, accidentally, or due to neglect, including serious crimes, acts of violence (e.g., martial arts applications), or even acts of war. This potential is amplified by the ease of DIY modification; for example, a hobbyist could easily 'DIY mod it into a logging robot with a chainsaw, made in their garage,' transforming a benign platform into a dangerous tool. While not designed as weapons, these platforms can be purposefully modified through readily available programming and hardware adaptations:

Covert Surveillance & Tracking: Utilizing their advanced sensors and mobility for unauthorized data collection and tracking in public or private spaces, bypassing privacy safeguards.

Targeted Harassment/Intimidation: Employing their physical presence and AI-driven social cues to harass, intimidate, or create discomfort in targeted individuals, with the "emotional intelligence" upgrades ironically enhancing this capability.

Physical Attack: While the G1 is not a combat robot, its 35kg mass and ability to navigate complex terrain and interact physically could, with malicious reprogramming and integration of additional components, be used to push, trip, or impact individuals. This parallels the concern with 3D printers enabling "home printing" of crucial, unregulated assault rifle components.

Unsanctioned Delivery/Deployment: Used for delivering illicit payloads or deploying other harmful devices in public or private areas, mirroring concerns seen with modified commercial drones capable of carrying objects.

This test case scenario powerfully illustrates that the legal and ethical vacuum surrounding commercially available, highly integratable, and upgradeable semi-autonomous humanoid robots like the Unitree G1 is not a hypothetical concern; the warning comes too late—it is already yesterday. It is an immediate and present challenge that demands urgent, comprehensive regulatory action to manage the spectrum from friendly assistance to unforeseen hazard, which is now squarely within the public's reach.

The Role of Architected Meaning for Enhanced Safety and Control:

Crucially, the inherent risks discussed above can be significantly mitigated through a deliberate architectural approach to AI design, moving beyond the limitations of general-purpose LLMs and traditional machine learning. As explored in Berend Watchus's recent work (e.g., Watchus, 2024, 2025) on the "synthetic insula" and "architectures of meaning," by explicitly engineering AI with "meaningful interfaces" and a "synthetic insula," it is possible to create systems that actively understand why certain situations are dangerous or undesirable, rather than simply recognizing patterns. For instance, a conventional self-driving car, trained on typical road data, might perform excellently in standard conditions. However, if it encounters a carnival parade where people are in elaborate costumes, or vehicles are disguised as space or ocean ships, a purely pattern-matching AI might "not compute" or fail to "register" these entities as humans or standard vehicles, leading to confusion or unsafe behavior. In contrast, a humanoid robot equipped with overall insight and meaning, stemming from a synthetic insula, would understand the context of a carnival, recognize that people in costumes are still humans, and interpret unusual shapes as artistic expressions rather than threats or malfunctions. This deeper understanding of why something is 'bad' (dangerous) allows it to maintain a human-centered checklist of meaning, always checking situations and actions with this ethical framework, thus making it significantly smarter and safer than a basic beginner software. This represents a significant shift towards truly controllable and ethically aligned AI.

Conclusion: Navigating the Inevitable Future

The exploration of the Unitree G1 scenario has revealed a compelling and immediate truth: the future of advanced robotics is not a distant concern, but a present reality actively unfolding within public accessibility. The availability of highly capable, modifiable platforms like the G1, purchasable via mail order and readily integratable with sophisticated AI, fundamentally shifts the dual-use dilemma from theoretical debate to urgent practical challenge. This paper has illuminated the inherent tension between the immense potential for benevolent applications and the demonstrable risks of misuse and unforeseen accidents that arise from an existing regulatory vacuum.

Our technical testing scenario in Part 1 showcased how easily advanced functionalities, such as self-recognition and empathetic response, can be achieved by enthusiasts through the integration of readily available computational resources

like LLMs and smartphone applications. This democratic access to powerful capabilities simultaneously opens doors to unprecedented innovation and benevolent uses, such as personalized assistance and humanitarian aid, while starkly outlining the potential for complex, unpredicted accidents and deliberate repurposing for harm.

Part 2 underscored the critical inadequacy of current legal frameworks in addressing these novel challenges. As demonstrated by the hypothetical lawsuit, existing laws often fall short when confronted with the nuanced realities of semi-autonomous, modifiable robots operating in public spaces. The emerging EU AI Act signals a crucial step forward, yet its full implementation and proactive enforcement remain paramount to prevent reactive and ambiguous legal responses to incidents that are already possible today. The urgent need for comprehensive regulatory action, encompassing clear safety protocols, liability frameworks, and ethical guidelines, is undeniable.

Crucially, this paper emphasizes that the path forward does not lie in halting technological progress, but in architecting it responsibly. The concepts of "meaning," "interface," and "synthetic insula," building upon foundational work in perception theory and self-awareness in AI, offer a vital blueprint. By designing AI systems with explicitly architected meaning, internal pseudo-affective states, and controlled feedback loops, we can cultivate systems that understand and prioritize ethical directives and human well-being, without aspiring to genuine sentience or independent goals. This approach ensures that AI's decision-making is rooted in a human-centered checklist of meaning, allowing it to navigate complex, unpredictable environments safely and effectively.

The Unitree G1 serves as a stark, immediate example: a commercially available, modifiable platform that embodies the spectrum from benevolent aid to unforeseen hazard. Proactive regulatory measures, coupled with a deep commitment to ethically aligned AI design informed by cutting-edge research into artificial meaning and control, are not options for a distant future. They are immediate imperatives for safeguarding society and ensuring that the transformative potential of advanced robotics serves humanity responsibly and equitably. The time to act and shape this future is now.

References

1. AZoRobotics. (2025, January 15). Navigating the EU AI Act: Implications for

the Robotics Industry. Retrieved July 1, 2025, from https://www.azorobotics.com/Article.aspx?ArticleID=736

2. ASU News. (2024, July 15). The ethical costs of advances in AI. Retrieved July 1, 2025, from https://news.asu.edu/20240715-law-journalism-and-politics-ethical-costs-advances-ai

3. Center for AI Safety. (n.d.). AI Risks that Could Lead to Catastrophe. Retrieved July 1, 2025, from https://safe.ai/ai-risk

4. European Commission. (n.d.). European approach to artificial intelligence. Retrieved July 1, 2025, from https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence

5. European Parliament. (n.d.). Regulation (EU) 2016/679 (General Data Protection Regulation). (General knowledge, widely accessible)

6. European Union. (n.d.). AI Act | Shaping Europe's digital future. Retrieved July 1, 2025, from https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

7. IEEE SA. (n.d.). Autonomous and Intelligent Systems (AIS) Standards. Retrieved July 1, 2025, from https://standards.ieee.org/initiatives/autonomous-intelligence-systems/standards/

8. IEEE Spectrum. (2022, January). AI's Real Worst-Case Scenarios. Retrieved July 1, 2025, from https://read.nxtbook.com/ieee/spectrum/spectrum_na_january_2022/ai_s_real_worst_case_scenario.html

9. National Robotics Programme. (n.d.). Standards. Retrieved July 1, 2025, from https://www.nrp.gov.sg/engineering/standards/

10. Number Analytics. (2025, June 11). Robot Emotion in AI Robotics. Retrieved July 1, 2025, from https://www.numberanalytics.com/blog/ultimate-guide-robot-emotion-ai-robotics

11. Number Analytics. (2025, June 23). Emotional Intelligence in Robotics. Retrieved July 1, 2025, from https://www.numberanalytics.com/blog/emotional-intelligence-robotics

12. OSTI.GOV. (n.d.). Evolution of the IEEE P7009 Standard: Towards Fail-Safe Design of Autonomous Systems. Retrieved July 1, 2025, from https://www.osti.gov/servlets/purl/1860583

13. PNAS. (2025, April 17). Existential risk narratives about AI do not distract from its immediate harms. Retrieved July 1, 2025, from https://www.pnas.org/doi/10.1073/pnas.2419055122

14. Qviro Blog. (2025, March 4). Unitree G1: The Affordable Humanoid Robot. Retrieved July 1, 2025, from https://qviro.com/blog/unitree-g1-the-affordable-humanoid-robot/

15. ResearchGate. (2024, July 24). AI-Driven Emotional Intelligence for Enhanced Human-Robot Interaction. Retrieved July 1, 2025, from https://www.researchgate.net/publication/382523407_AI-Driven_Emotional_Intelligence_for_Enhanced_Human-_Robot_Interaction

16. RobotShop. (n.d.). Unitree G1 Humanoid Robot (EU). Retrieved July 1, 2025, from https://eu.robotshop.com/products/unitree-g1-humanoid-robot-eu

17. SiliconANGLE. (2025, March 8). AI's existential risks: Separating hype from reality. Retrieved July 1, 2025, from https://siliconangle.com/2025/03/08/ais-existential-risks-separating-hype-reality/

18. The Wassenaar Arrangement. (n.d.). List of Dual-Use Goods and Technologies and Munitions List. Retrieved July 1, 2025, from https://www.wassenaar.org/app/uploads/2023/12/List-of-Dual-Use-Goods-and-Technologies-Munitions-List-2023-1.pdf

19. Watchus, B. (2024, Nov 14). Advanced Predictive Modeling of Physical Trajectories and Cascading Events, Dual-State Feedback and Synthetic Insula. https://doi.org/10.20944/preprints202411.1025.v1

20. Watchus, B. (2024, Nov 15). Self-Identification in AI: ChatGPT's Current Capability for Mirror Image Recognition. https://doi.org/10.20944/preprints202411.1112.v1

21. Watchus, B. (2024, Nov 12). Simulating Self-Awareness: Dual Embodiment, Mirror Testing, and Emotional Feedback in AI Research https://doi.org/10.20944/preprints202411.0839.v1

22. Watchus, B. (2024, Nov 12). The Unified Model of Consciousness: Interface and Feedback Loop as the Core of Sentience. https://doi.org/10.20944/preprints202411.0727.v1

23. Watchus, B. (2024, Nov 11). Towards Self-Aware AI: Embodiment, Feedback Loops, and the Role of the Insula in Consciousness. https://doi.org/10.20944/preprints202411.0661.v1

24. Watchus, B. (2025, June 25). The Architectures of Meaning: Integrating Hoffman's Perception Theory with Synthetic Ethical Embodiment in AI. https://doi.org/10.20944/preprints202506.2025.v1

25. Wikipedia. (n.d.). Existential risk from artificial intelligence. Retrieved July 1, 2025, from https://en.wikipedia.org/wiki/Existential_risk_from_artificial_intelligence